Lenovo

Balanced Memory Configurations with Second-Generation Intel Xeon Scalable Processors

Demonstrates balanced memory guidelines for second-generation Intel Xeon SP processors

Compares the performance of balanced and unbalanced memory configurations

Explains memory interleaving and its importance

Provides tips on how to balance memory and maximize performance

Dan Colglazier

Joseph Jakubowski

Jamal Ayoubi

LENOVO PRESS

Abstract

Configuring a server with balanced memory is important for maximizing its memory bandwidth and overall performance. Lenovo ThinkSystem servers running second-generation Intel Xeon Scalable processors have six memory channels per processor and up to two DIMMs per channel, so it is important to understand what is considered a balanced configuration and what is not.

This paper defines three balanced memory guidelines that will guide you to select a balanced memory configuration. Balanced and unbalanced memory configurations are presented along with their relative measured memory bandwidths to show the effect of unbalanced memory. Suggestions are also provided on how to produce balanced memory configurations.

This paper is for customers and for business partners and sellers wishing to understand how to maximize the performance of Lenovo ThinkSystem servers with second-generation Intel Xeon Scalable processors.

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges.

See a list of our most recent publications at the Lenovo Press web site:

http://lenovopress.com

Do you have the latest version? We update our papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest if needed. While you're there, you can also sign up to get notified via email whenever we make an update.

Contents

Introduction

The memory subsystem is a key component of the Intel x86 server architecture and it can greatly affect overall server performance. When properly configured, the memory subsystem can deliver extremely high memory bandwidth and low memory access latency. When the memory subsystem is incorrectly configured, the memory bandwidth available to the server can become limited and overall server performance can be severely reduced.

This brief explains the concept of balanced memory configurations that yield the highest possible memory bandwidth from the Intel Xeon Scalable processors that are used in Lenovo® ThinkSystem™ servers. By increasing the number of populated DIMMs from one to twelve, examples of balanced and unbalanced memory configurations are shown to illustrate their effect on memory subsystem performance.

This brief specifically covers the second-generation Intel Xeon Scalable processors (Xeon SP) that were announced in April 2019. For first-generation processors, see the paper *Intel Xeon Scalable Family Balanced Memory Configurations*, available from:

https://lenovopress.com/lp0742

Memory interleaving

Access to the information stored on DIMMs is controlled by the memory controllers that are integrated within the processor. With Intel Xeon Scalable processors, two memory controllers are present. Each memory controller is attached to three memory channels that are connected to the physical slot connectors that interface to the DIMMs.

Figure 1 illustrates how a Xeon SP processor's memory controllers are connected to memory DIMM slots.

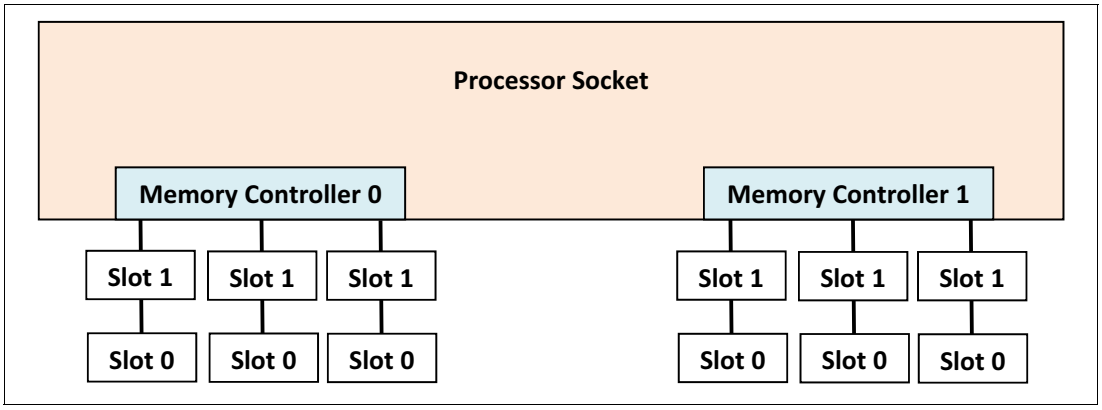

Figure 1 Intel Xeon SP processor with two memory controllers, six memory channels and twelve memory DIMM slots

The Xeon SP processors optimizes memory accesses by creating interleave sets across the memory controllers and memory channels. For example, if two memory channels have the same total memory capacity, a 2-way interleave set is created across the memory channels. Interleaving enables higher memory bandwidth by spreading contiguous memory accesses across both memory channels rather than sending all memory accesses to one memory channel.

If DIMMS are populated on the memory channels such that they have different total memory capacities, the memory controller has to create multiple interleave sets. Some interleave sets could have fewer DIMMs. Managing multiple interleave sets creates overhead for the memory controllers, which can reduce memory bandwidth.

In addition, the performance of a specific memory access depends on which memory region is being accessed and how many DIMMs comprise the interleave set. Contiguous memory accesses to a memory region with fewer DIMMs in the interleave set will have lower performance compared to accesses to a memory region with more DIMMs in the interleave set.

Figure 2 illustrates a 4-channel interleave set on a Scalable processor that results from populating identical DIMMs on two memory channels on each memory controller. This 4-channel set interleaves across the memory controllers and between the memory channels on each memory controller. Consecutive addresses alternate between the memory controllers with every fourth address going to each memory channel.

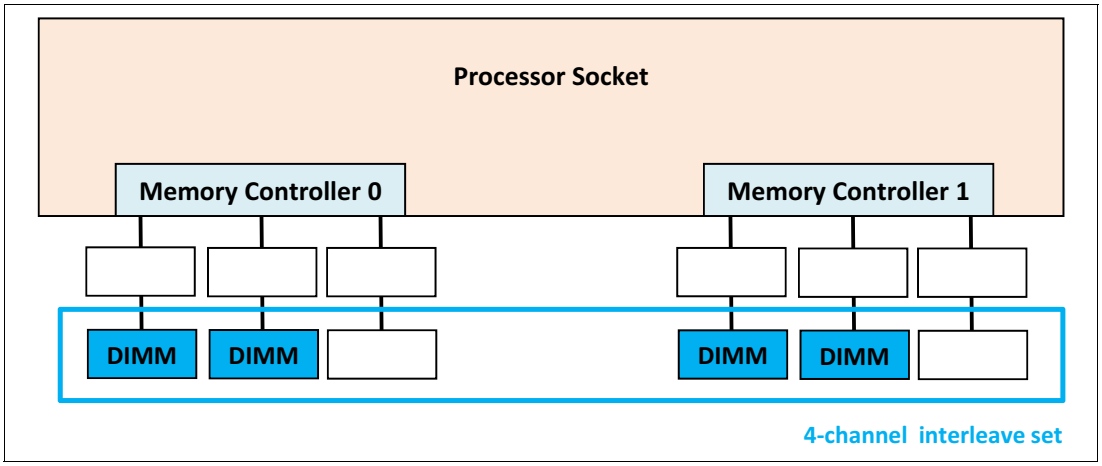

Figure 2 4-channel interleave set across memory controllers and between memory channels

Within a memory channel, a second level of interleaving called *memory rank interleaving* can occur. A *memory rank* is a block of data created from the memory chips on a memory DIMM. A memory rank is typically 64 bits wide. If ECC is supported, an additional 8 bits are added for a total of 72 bits. A DIMM may contain multiple memory ranks with one, two and four rank DIMMs being the most common.

Memory rank interleaving generally improves memory performance as the total number of ranks on a memory channel increases, but only up to a point. The Intel architecture is optimized for two to four memory ranks per memory channel. Beyond four ranks per memory channel, performance can slightly degrade due to electrical turnaround time on the memory channel when the memory controller switches between memory ranks.

Balanced memory configurations

Balanced memory configurations enable optimal interleaving, which maximizes memory bandwidth. Optimal memory bandwidth occurs when all the populated memory channels have the same total memory capacity and total number of ranks. Memory bandwidth is optimal when all memory controllers on the same physical processor socket are identically configured. System level memory bandwidth is optimal when each physical processor socket has the same physical memory capacity.

The basic guidelines for a balanced memory subsystem are therefore as follows:

1. All populated memory channels should have the same total memory capacity and the same total number of ranks
2. All memory controllers on a processor socket should have the same configuration of DIMMs
3. All processor sockets on the same physical server should have the same configuration of DIMMs

Tip: We will refer to the above guidelines as *Balanced Memory Guidelines 1, 2 and 3* throughout this brief.

About the tests

STREAM Triad is a simple, synthetic benchmark designed to measure sustainable memory bandwidth. Its intent is to measure the best memory bandwidth available. STREAM Triad will be used to measure the sustained memory bandwidth of various memory configurations to see the effect of suboptimal memory configurations on memory bandwidth.

For more information about STREAM Triad, see the following web page:

http://www.cs.virginia.edu/stream/

Memory topology

A Xeon SP processor has two memory controllers. Each memory controller has three memory channels, and each memory channel supports one or two DIMM slots. To illustrate various memory topologies for a processor with two memory controllers, different memory configurations will be designated as `A:B:C,D:E:F` where each letter indicates the number of DIMMs populated on each memory channel.

- A refers to Memory Channel 0 on Memory Controller 0
- B refers to Memory Channel 1 on Memory Controller 0
- C refers to Memory Channel 2 on Memory Controller 0
- D refers to Memory Channel 0 on Memory Controller 1
- E refers to Memory Channel 1 on Memory Controller 1
- F refers to Memory Channel 2 on Memory Controller 1

As an example, a `2:2:2,1:1:1` memory configuration has:

- 2 DIMMs on Memory Channels 0, 1, and 2 on Memory Controller 0
- 1 DIMM on Memory Channels 0, 1, and 2 on Memory Controller 1.

Applying the balanced memory configuration guidelines

All tests in this paper are using server configurations with second-generation Intel Xeon Scalable processors.

We will start with the assumption that Balanced Memory Guideline 3 (described in “Balanced memory configurations” on page 4) is followed: all processor sockets on the same physical

server have the same configuration of DIMMs. Therefore, we only have to look at one processor to describe each memory configuration.

All DIMMs used are 32 GB dual rank RDIMMs. The number of these DIMMs used will be increased from one to twelve to see the effect on memory bandwidth. For each memory configuration we will determine which balanced memory guidelines are followed, and the number and type of interleave sets will be shown. Any recommendations for improving the performance of the memory configuration will also be pointed out.

Installation sequence: When installing DIMMs, follow the DIMM installation sequence for that particular server. The configurations shown in this brief did not always follow the sequences for the server they were measured on because a number of these configurations were put together just for demonstration purposes.

Configuration of 1 DIMM - unbalanced

We will start with one 32GB dual-rank DIMM, which yields the 1:0:0,0:0:0 memory configuration shown in Figure 3.

Balanced memory guideline 1 is followed with only one populated memory channel. Balanced memory guideline 2 is not followed as only one memory controller is populated. This is not a balanced memory configuration.

A single 1-channel interleave set is formed. Having only one memory channel populated with memory greatly reduces the memory bandwidth of this configuration, which was measured at 18% or about one sixth of the full potential memory bandwidth.

The best way to increase the memory bandwidth of this configuration is by using more DIMMs. Two 16 GB dual-rank RDIMMs would provide the same memory capacity while nearly doubling the memory bandwidth.

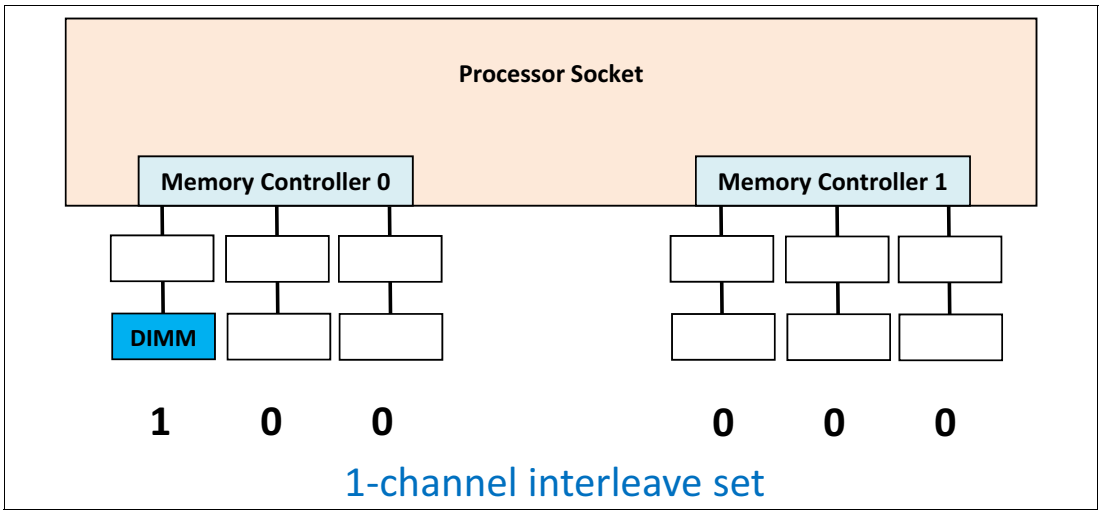

Figure 3 1:0:0,0:0:0 memory configuration (STREAM Triad relative memory bandwidth = 18%)

Configuration of 2 DIMMs - balanced if installed correctly

Two DIMMs can be configured in two different ways, one will be balanced but the other will not.

The first we will look at is the 1:0:0,1:0:0 memory configuration shown in Figure 4 on page 7. This memory configuration follows balanced memory guideline 1, because both populated memory channels have the same memory capacity. It also follows balanced memory guideline 2 with the same configuration on each memory controller. This is a balanced memory configuration.

A single 2-channel interleave set is formed across the memory controllers. Only two memory channels are populated with memory, which greatly reduces the memory bandwidth of this memory configuration to about one-third of the full potential memory bandwidth. It was measured at 35%.

The best way to increase the memory bandwidth of this configuration is by using more DIMMs. Four 16 GB RDIMMs would provide the same memory capacity while nearly doubling memory bandwidth.

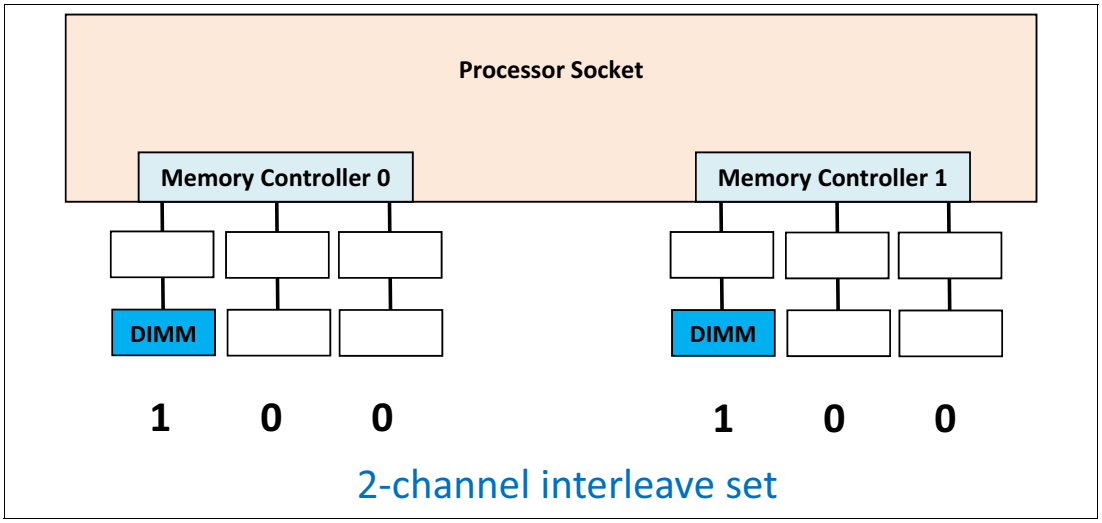

Figure 4 1:0:0,1:0:0 memory configuration (STREAM Triad relative memory bandwidth = 35%)

The second way to arrange two DIMMs is to attach both of them to the same memory controller as in the 1:1:0,0:0:0 memory configuration shown in Figure 5 on page 8. This memory configuration does follow balanced memory guideline 1 with both populated memory channels having the same memory capacity. It does not follow balanced memory guideline 2 having different configurations on the memory controllers. This is not a balanced memory configuration.

A single 2-channel interleave set is formed on the one populated memory controller. Only two memory channels are populated with memory, greatly reducing the bandwidth of this memory configuration to about one third of the full potential memory bandwidth. It was measured at 35% showing interleaving between two channels on a memory controller provides the same memory bandwidth as interleaving between memory controllers.

The best way to increase the memory bandwidth of this configuration is by using more DIMMs.

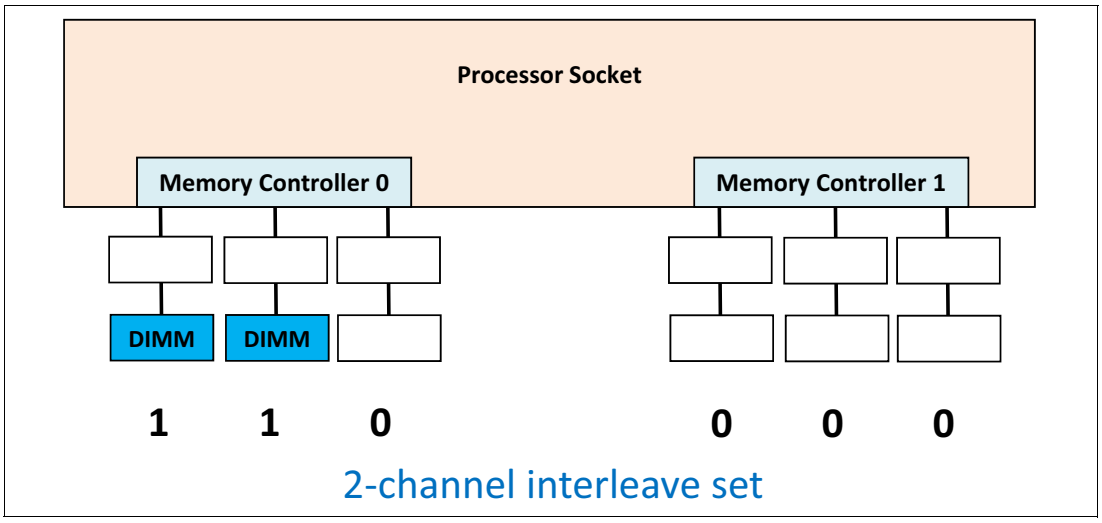

Figure 5 1:1:0,0:0:0 memory configuration (STREAM Triad relative memory bandwidth = 35%)

Configuration of 3 DIMMs - unbalanced

Three DIMMs can all be attached to the same memory controller or spread between two memory controllers.

Memory bandwidth is better for the 1:1:1,0:0:0 memory configuration shown in Figure 6. This configuration does follow balanced memory guideline 1 as the three populated memory channels have the same memory capacity. It does not follow balanced memory guideline 2 having different configurations on each memory controller. It is not a balanced memory configuration.

A single 3-channel interleave set is formed on the one populated memory controller. Only half of the memory channels are populated with memory, reducing the memory bandwidth of this memory configuration to about half of the full potential memory bandwidth. It was measured at 51%. The best way to increase the memory bandwidth of this memory configuration is by populating all memory channels by using six 16GB DIMMs.

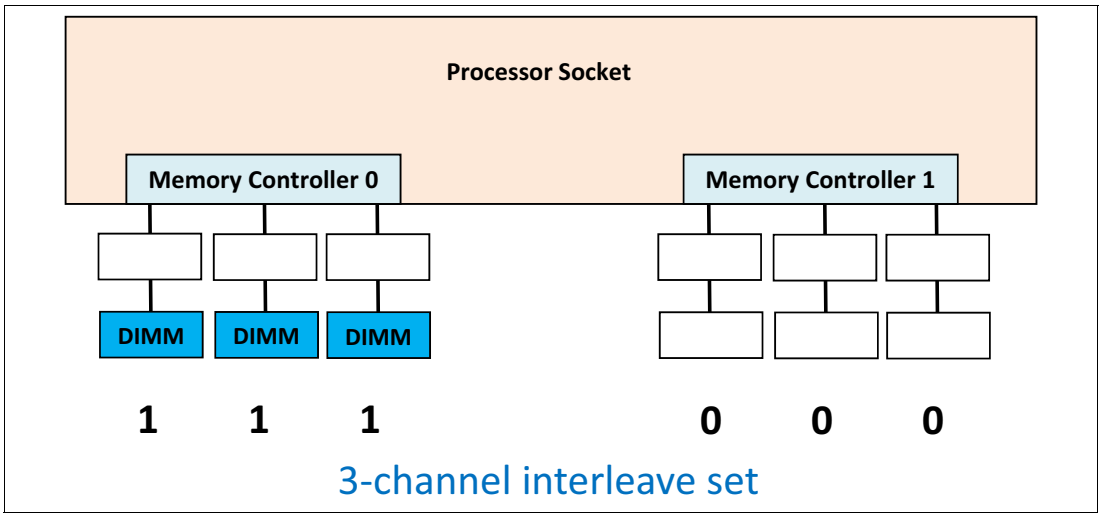

Figure 6 1:1:1,0:0:0 memory configuration (STREAM Triad relative memory bandwidth = 51%)

Spreading three DIMMs across two memory controllers greatly reduces memory bandwidth compared to being on the same memory controller. This 1:1:0,1:0:0 memory configuration shown in Figure 7 does follow balanced memory guideline 1 but not 2. It is not a balanced memory configuration.

Two interleave sets are formed for this memory configuration: one 2-channel interleave set across the memory controllers and one 1-channel interleave set with the remaining memory. Having more than one interleave set greatly reduces memory bandwidth. This memory configuration was measured at 18% of the full potential memory bandwidth, despite populating half of the memory channels. This demonstrates the importance of having a single interleave set for optimal memory bandwidth.

The best ways to increase the memory bandwidth of this configuration is by forming only one interleave set by attaching all three DIMMs to the same memory controller and/or by populating all memory channels by using six 16GB memory DIMMs.

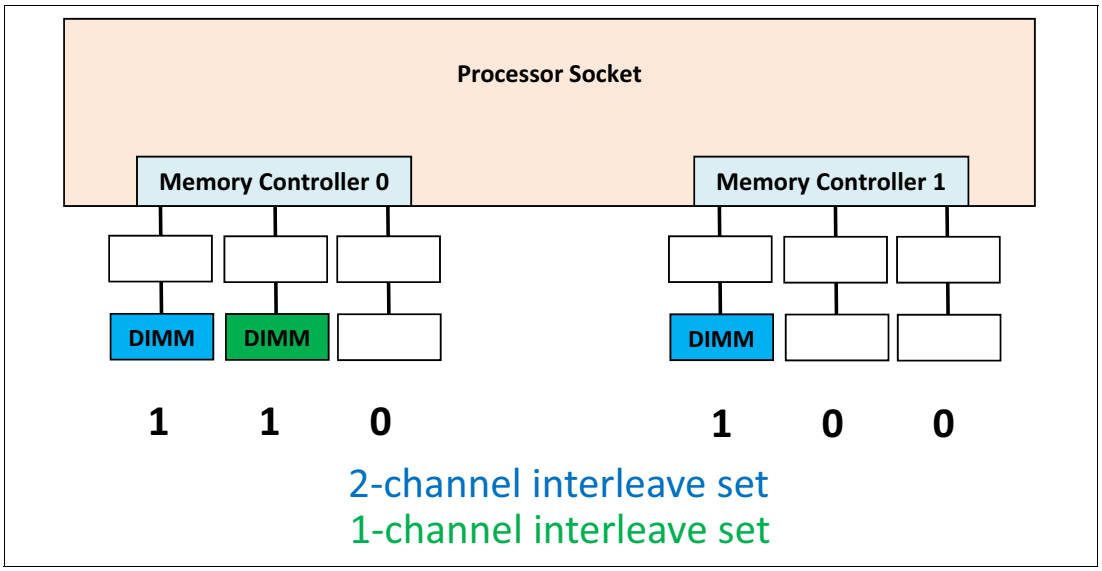

Figure 7 1:1:0,1:0:0 memory configuration (STREAM Triad relative memory bandwidth = 18%)

Configuration of 4 DIMMs - balanced if installed correctly

Four DIMMs can be populated in the 1:1:0,1:1:0 memory configuration shown in Figure 8 on page 10. This memory configuration follows balanced memory guideline 1, because the four populated memory channels have the same memory capacity. It also follows balanced memory guideline 2, by having the same memory configuration on each memory controller. This is a balanced memory configuration.

A single 4-channel interleave set is formed across the memory controllers. Only four of the memory channels are populated with memory, reducing the memory bandwidth of this configuration to about two thirds of the full potential memory bandwidth. It was measured at 66%.

The best way to increase the memory bandwidth of this memory configuration is by populating all memory channels by using more DIMMs.

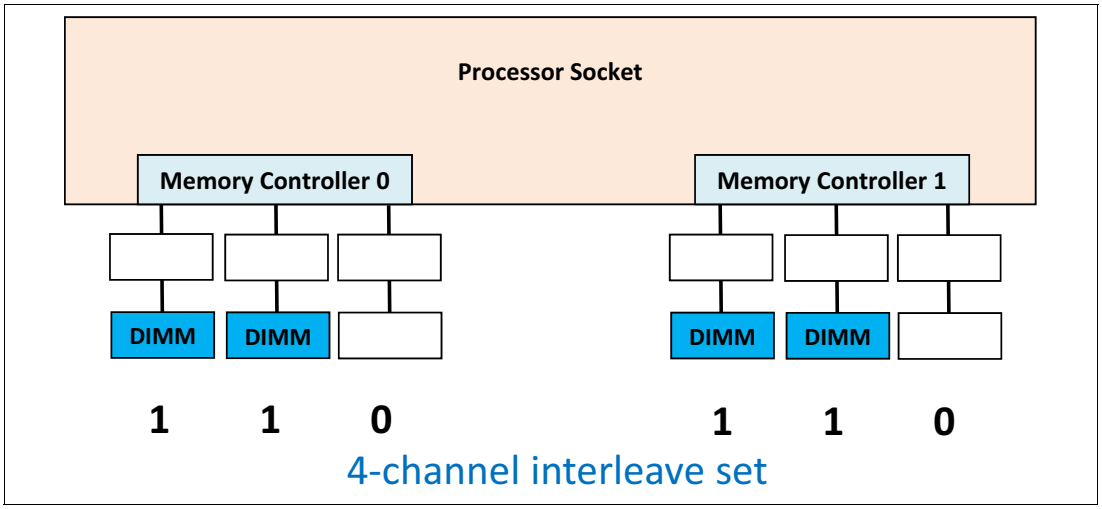

Figure 8 1:1:0,1:1:0 memory configuration (STREAM Triad relative memory bandwidth = 66%)

Four DIMMs can also be populated in the 1:1:1,1:0:0 memory configuration shown in Figure 9. This memory configuration follows balanced memory guideline 1 but not 2. It is not a balanced memory configuration.

Two 2-channel interleave sets are formed, one across the memory controllers and one with the remaining memory on a single memory controller. As seen before, more than one interleave set is detrimental to memory bandwidth. This memory configuration was measured at 35% which is about half of the bandwidth of four DIMMs in one interleave set.

The best way to increase the memory bandwidth of this memory configuration is by moving one memory DIMM to reduce the number of interleave sets from two to one.

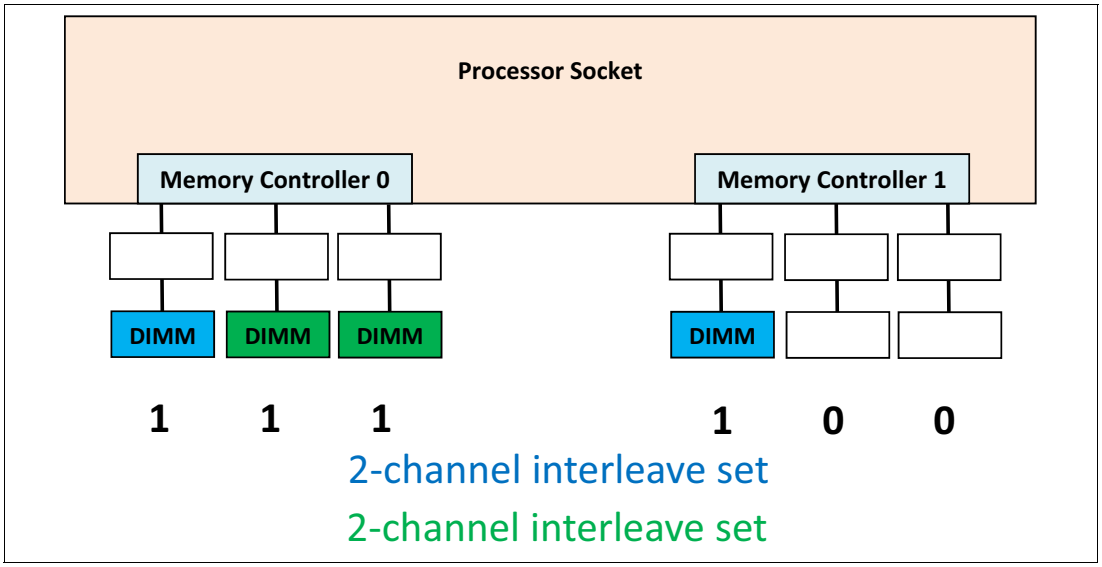

Figure 9 1:1:1,1:0:0 memory configuration (STREAM Triad relative memory bandwidth = 35%)

Configuration of 5 DIMMs - unbalanced

A configuration of five DIMMs is best populated in the 1:1:1,1:1:0 memory configuration shown in Figure 10 on page 11. While this memory configuration does follow balanced memory guideline 1 with the five populated memory channels having the same memory

capacity, it does not follow balanced memory guideline 2 with differing memory configurations on each memory controller. It is not a balanced memory configuration.

A 4-channel interleave set is formed across the memory controllers along with a 1-channel interleave set with the remaining memory. Having two interleave sets reduces the bandwidth of this memory configuration to a measured 18%.

The best way to increase the memory bandwidth of this memory configuration is by adding another DIMM, thereby populating all the memory channels and forming a single interleave set.

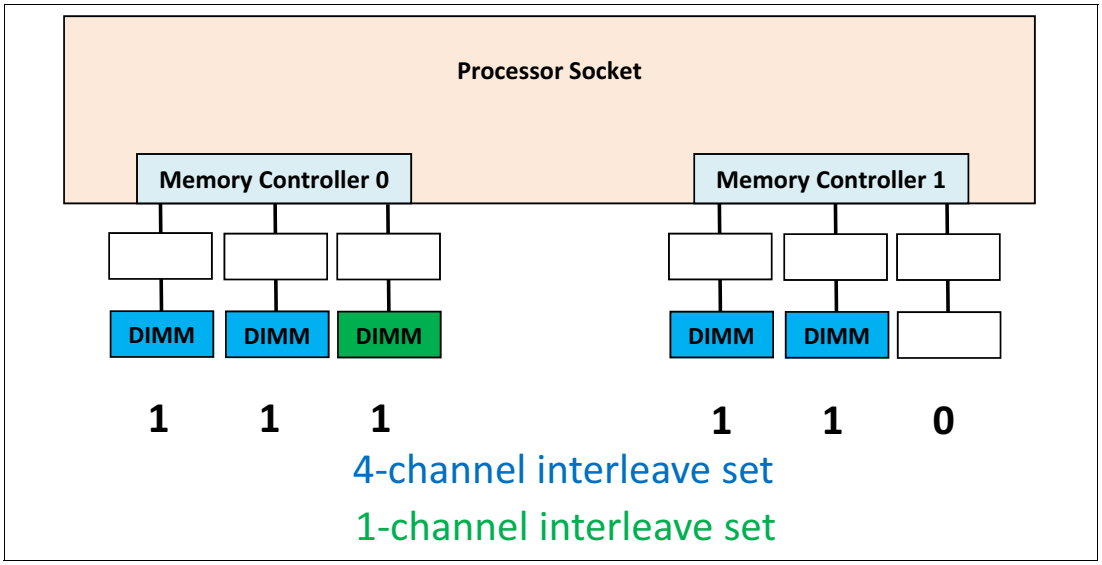

Figure 10 1:1:1,1:1:0 memory configuration (STREAM Triad relative memory bandwidth = 18%)

Configuration of 6 DIMMs - balanced

A balanced memory configuration is achieved with six DIMMs in the 1:1:1,1:1:1 memory configuration shown in Figure 11 on page 12. This configuration follows balanced memory guideline 1, because all six memory channels have the same memory capacity. It also follows balanced memory guideline 2, having the same configuration on each memory controller.

A single 6-channel interleave set is formed across the memory controllers. All of the memory channels are populated with memory maximizing the potential memory bandwidth which was measured at 95%.

Having only two ranks per memory channel slightly reduces this configuration's memory bandwidth. Having four ranks per memory channel would increase memory bandwidth by 5% which is the best way to increase the memory performance of this memory configuration.

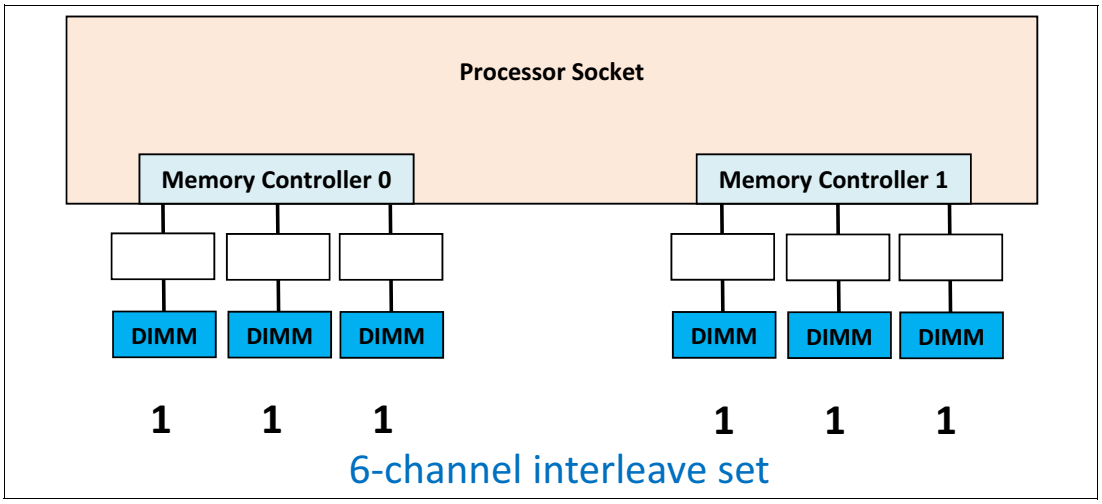

Figure 11 1:1:1,1:1:1 memory configuration (STREAM Triad relative memory bandwidth = 95%)

Configuration of 7 DIMMs - unbalanced

Seven DIMMs can be populated in the 2:1:1,1:1:1 memory configuration shown in Figure 12. This memory configuration does not follow balanced memory guideline 1 as one memory channel has twice the memory capacity of the other memory channels. It also does not follow balanced memory guideline 2 with differing configurations on each memory controller. This is not a balanced memory configuration.

A 6-channel interleave is formed across the memory controllers along with a 1-channel interleave set with the remaining memory. Having two interleave sets reduces the bandwidth of this memory configuration to a measured 18%.

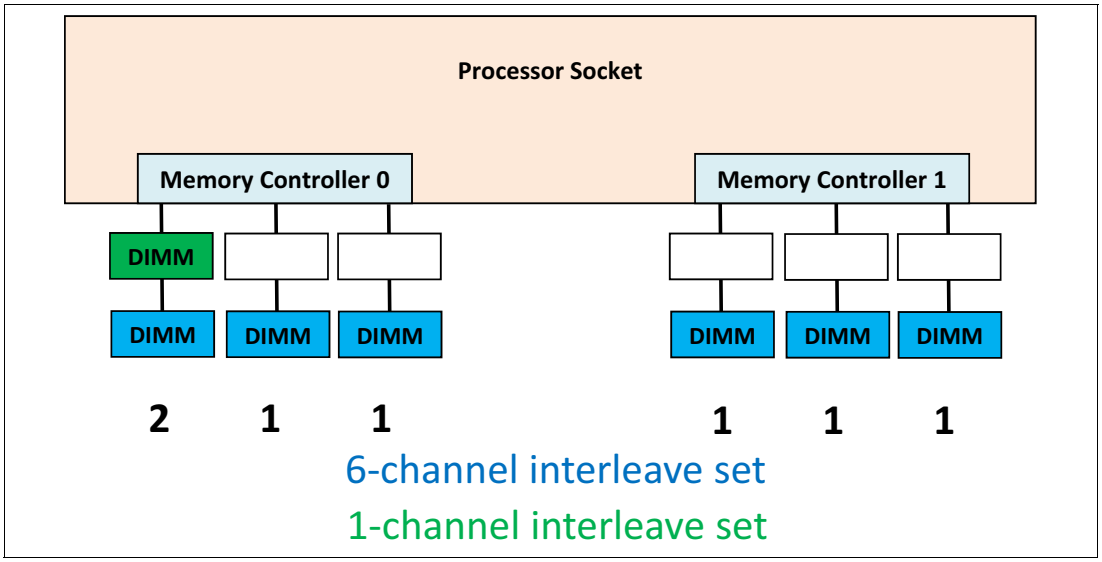

Figure 12 2:1:1,1:1:1 memory configuration (STREAM Triad relative memory bandwidth = 18%)

Populating seven DIMMs in the 2:2:0,1:1:1 memory configuration shown in Figure 13 on page 13 does not help memory bandwidth. It does not follow balanced memory guidelines 1 and 2, not all memory channels are populated, and three interleave sets are formed. This memory configuration was measured at 18%. It is best not to use seven DIMMs if memory bandwidth is important.

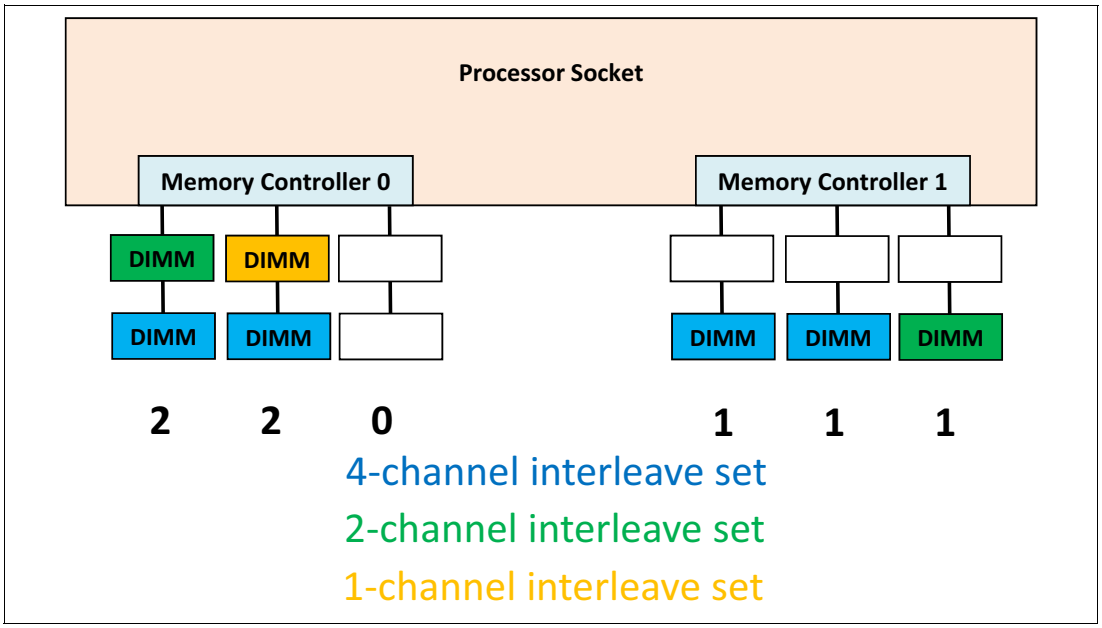

Figure 13 2:2:0,1:1:1 memory configuration (STREAM Triad relative memory bandwidth = 18%)

Configuration of 8 DIMMs - balanced if installed correctly

Eight DIMMs can be populated in the 2:2:0,2:2:0 memory configuration shown in Figure 14. This memory configuration follows balanced memory guideline 1 as the four populated memory channels have the same memory capacity. It also follows balanced memory guideline 2 having the same memory configuration on each memory controller. This is a balanced memory configuration.

A single 4-channel interleave set is formed across the memory controllers. Only four of the memory channels are populated with memory, reducing the memory bandwidth of this configuration to about two thirds of the full potential memory bandwidth. It was measured at 69%.

The best way to increase the memory bandwidth of this configuration is to populate more memory channels by using more DIMMs.

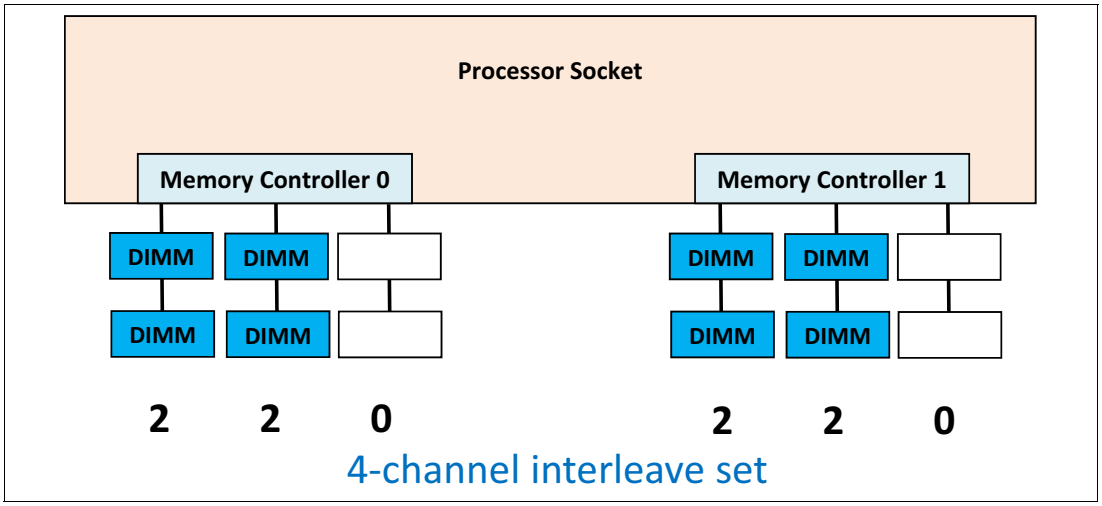

Figure 14 2:2:0,2:2:0 memory configuration (STREAM Triad relative memory bandwidth = 69%)

Eight DIMMs can also be populated in the 2:1:1,2:1:1 memory configuration shown in Figure 15. This memory configuration does not follow balanced memory guideline 1 as two of the memory channels have twice the memory capacity of the other memory channels. It does follow balanced memory guideline 2 with the same memory configuration on both memory controllers. This is not a balanced memory configuration.

A 6-channel interleave set and a 2-channel interleave set are formed across the memory controllers. As seen before, more than one interleave set is detrimental to memory bandwidth. This memory configuration was measured at 35% which is about half of the bandwidth with eight DIMMs in one interleave set.

The best way to increase the memory bandwidth of this memory configuration is by moving two DIMMs to reduce the number of interleave sets from two to one.

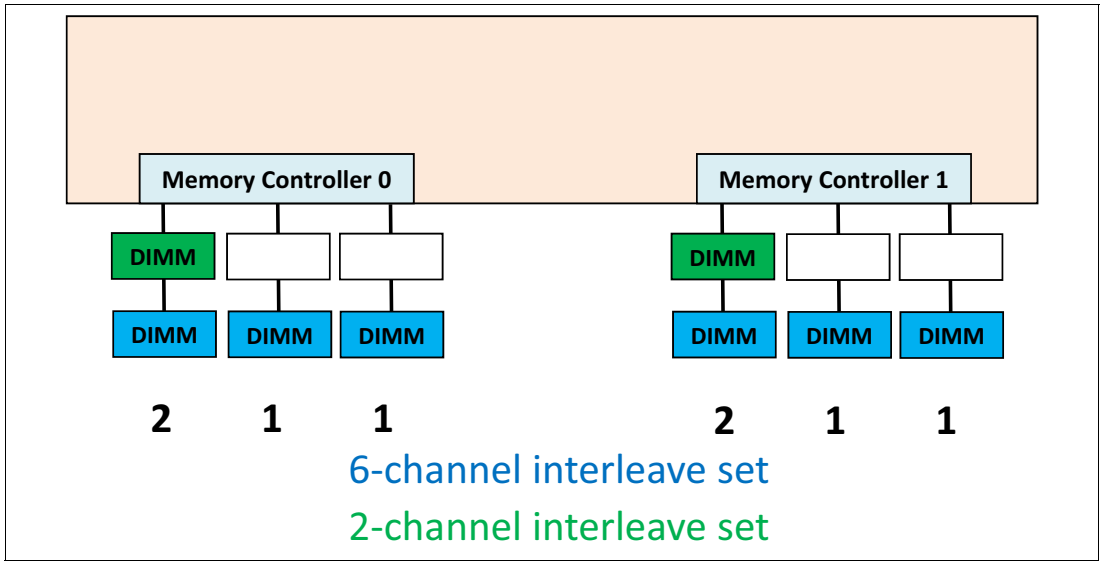

Figure 15 2:1:1,2:1:1 memory configuration (STREAM Triad relative memory bandwidth = 35%)

Configuration of 9 DIMMs - unbalanced

Populating nine DIMMs in two different memory configurations clearly demonstrates how having fewer interleave sets increases memory bandwidth. Neither one follows balanced memory guidelines 1 or 2. Neither one is a balanced memory configuration.

The 2:2:2,1:1:1 memory configuration is shown in Figure 16 on page 15. It forms two interleave sets: a 6-channel interleave set across the memory controllers and a 3-channel interleave set of the remaining memory. It was measured at 51% of the full potential memory bandwidth.

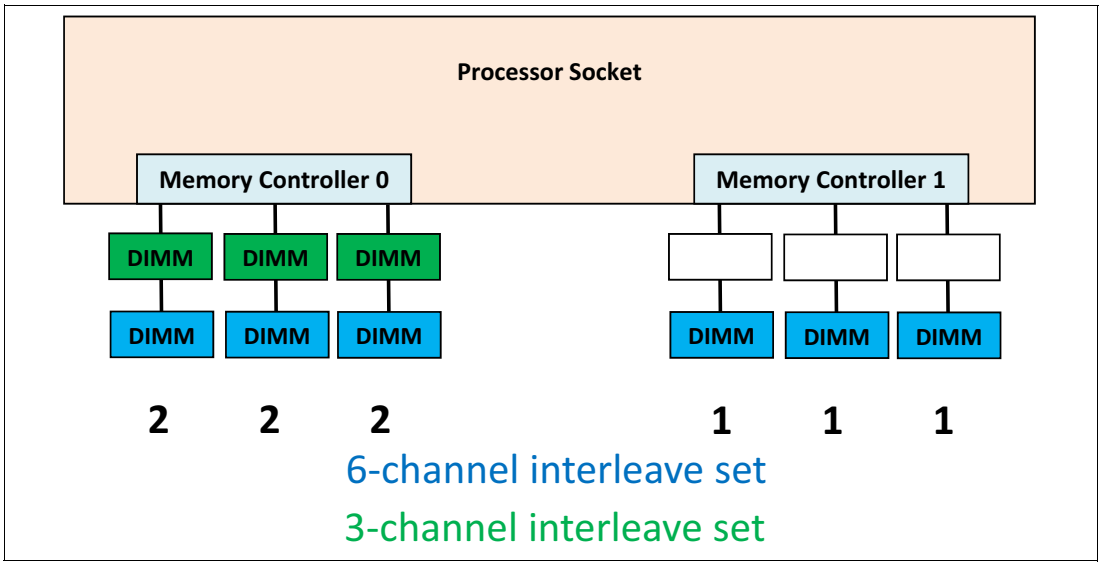

Figure 16 2:2:2,1:1:1 memory configuration (STREAM Triad relative memory bandwidth = 51%)

Moving one DIMM so that there is improved balance between the memory controllers produces the 2:2:1,2:1:1 memory configuration shown in Figure 17. Unfortunately, this results in three interleave sets being formed: a 6-channel and a 2-channel interleave set across the memory controllers and a 1-channel interleave set of the remaining memory. Three interleave sets is very detrimental to memory bandwidth. This memory configuration was measured at 18%. Memory bandwidth of a nine memory DIMM configuration can be improved by using three additional DIMMs.

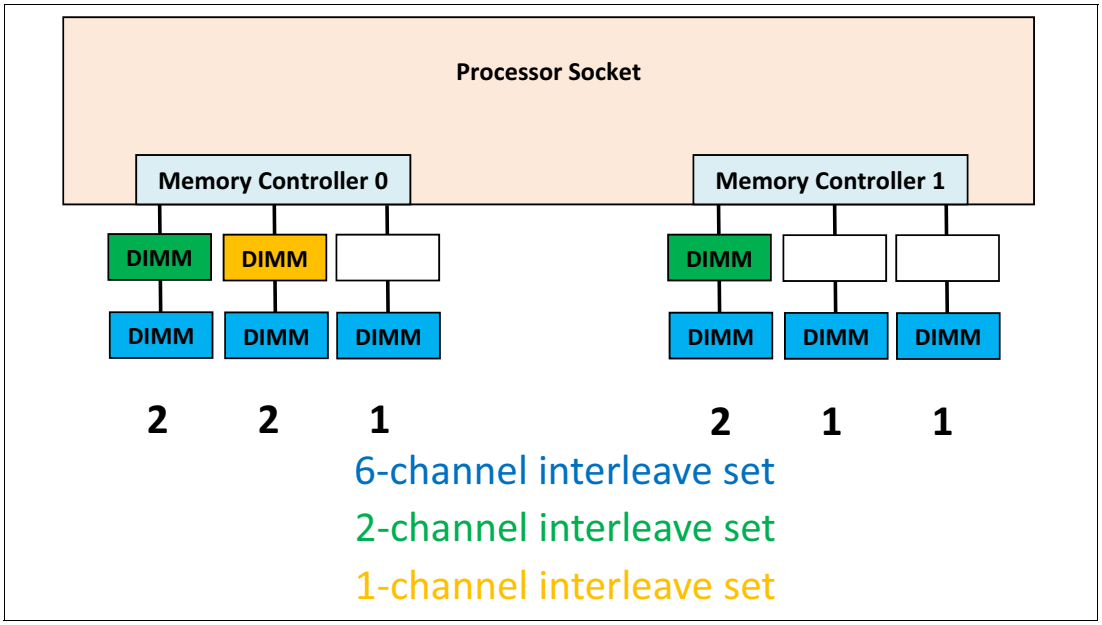

Figure 17 2:2:1,2:1:1 memory configuration (STREAM Triad relative memory bandwidth = 18%)

Configuration of 10 DIMMs - unbalanced

Populating ten DIMMs in the 2:2:1,2:2:1 memory configuration shown in Figure 18 on page 16 does not follow balanced memory guideline 1 because four memory chaniinels have twice the memory capacity as the other two. It does follow balanced memory guideline 2 with

the same memory configuration on both memory controllers. This is not a balanced memory configuration.

Two interleave sets are formed across the memory controllers: one 6-channel and one 4-channel interleave set. This combination works fairly well and was measured at 69% of full potential memory bandwidth.

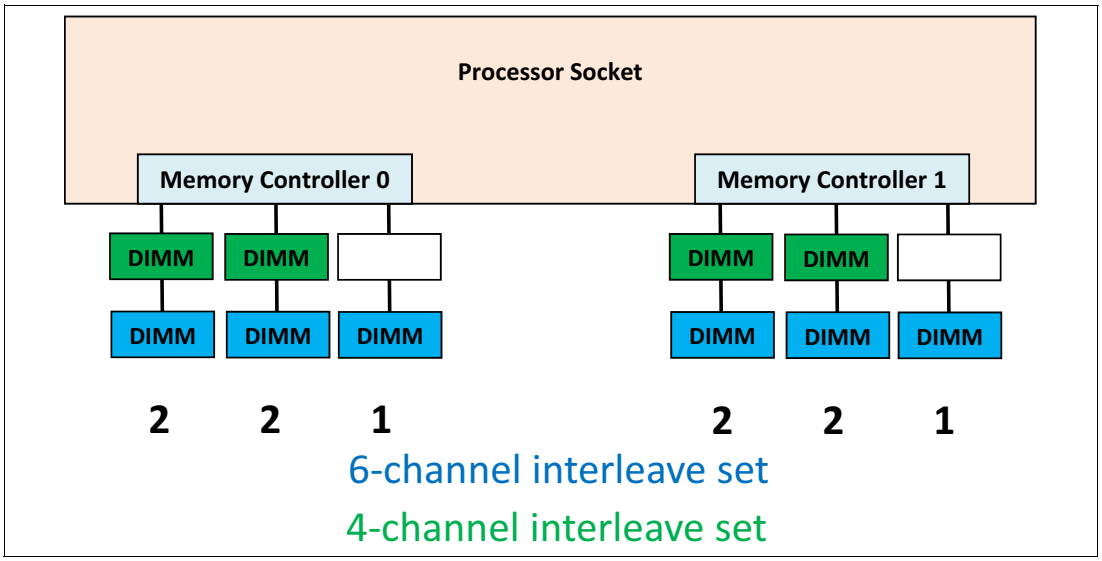

Figure 18 2:2:1,2:2:1 memory configuration (STREAM Triad relative memory bandwidth = 69%)

Moving one memory DIMM to create the 2:2:2,2:2:0 memory configuration shown in Figure 19 greatly reduced memory bandwidth down to 18% of full potential memory bandwidth. It does follow balanced memory guideline 1 but not 2 and is not a balanced memory configuration. This memory configuration forms a 4-channel interleave set across the memory controllers and a 1-channel interleave set with the remaining memory.

Ten-DIMM configurations can be improved by adding two additional DIMMs.

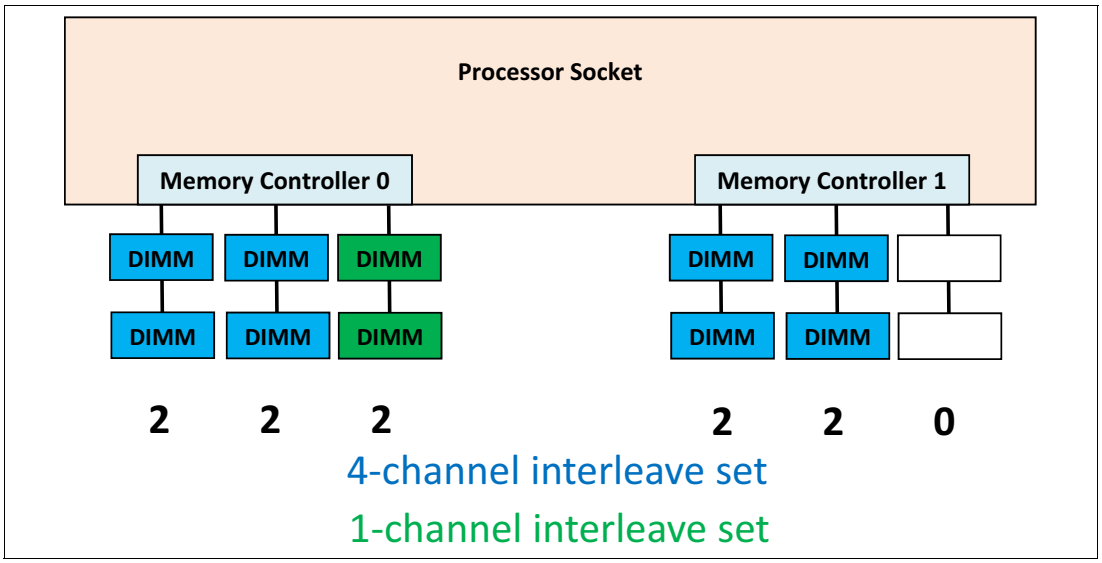

Figure 19 2:2:2,2:2:0 memory configuration (STREAM Triad relative memory bandwidth = 18%)

Configuration of 11 DIMMs - unbalanced

Populating eleven DIMMs in the 2:2:2,2:2:1 memory configuration shown in Figure 20 does not provide very good memory bandwidth. It does not follow either balanced memory guideline 1 or 2 and is not a balanced memory configuration. It forms three interleave sets: a 6-channel and a 4-channel interleave set across the memory controllers and a 1-channel interleave set with the remaining memory. The memory bandwidth was measured at 18% of the full potential memory bandwidth. Adding an additional memory DIMM will greatly improve memory bandwidth.

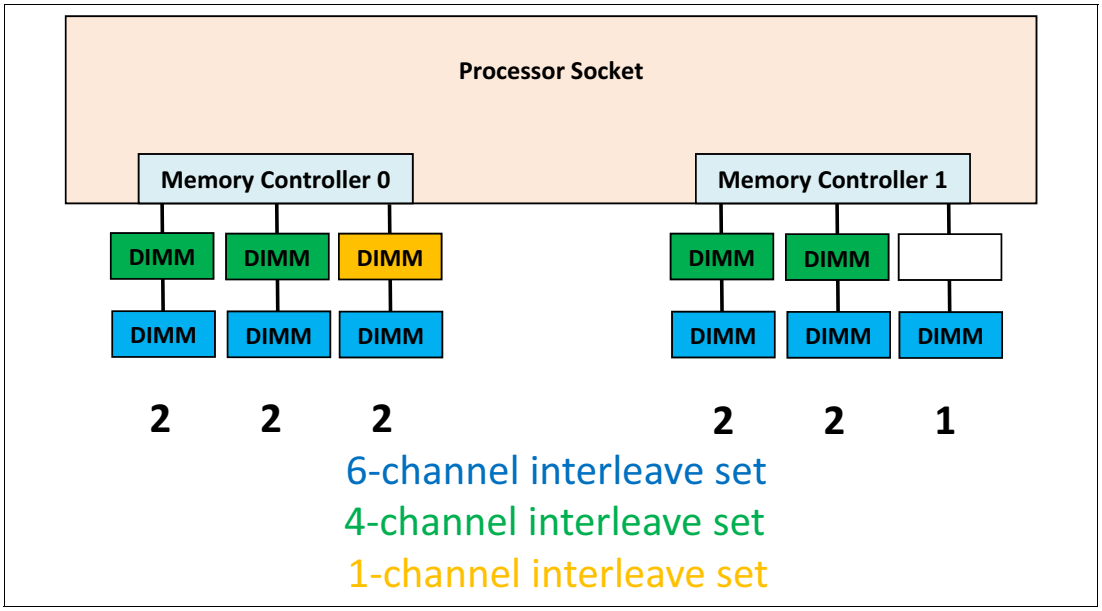

Figure 20 2:2:2,2:2:1 memory configuration (STREAM Triad relative memory bandwidth = 18%)

Configuration of 12 DIMMs - balanced and the best performance

Populating all twelve DIMM slots with memory provides all the potential memory bandwidth and was measured at 100%. This 2:2:2,2:2:2 memory configuration is shown in Figure 21 on page 18. This memory configuration is balanced because it follows both balanced memory guidelines 1 and 2 with the same memory capacity on each memory channel and the same memory configuration on each memory controller. Only one 6-channel interleave set across the memory controllers is formed. This is the ideal memory configuration for maximum performance with four ranks per memory channel.

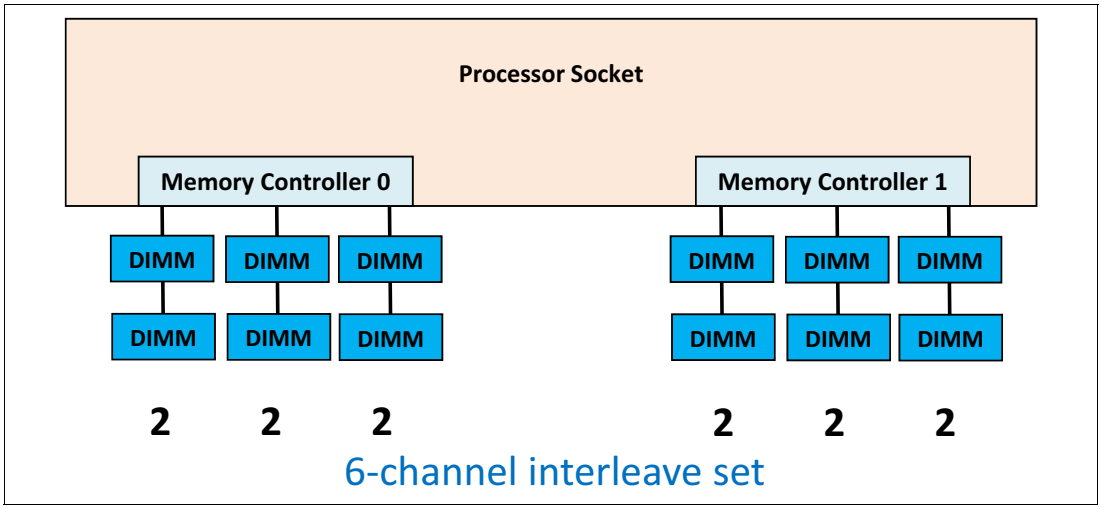

Figure 21 2:2:2,2:2:2 memory configuration (STREAM Triad relative memory bandwidth = 100%)

Summary of the performance results

Table 1 shows a summary of the relative memory bandwidth of all the memory configurations using second-generation Intel Xeon Scalable processors that are presented in this paper. It also shows the number of interleave sets formed for each and whether it is a balanced or unbalanced memory configuration.

When using the same DIMM, only memory configurations with 6 or 12 DIMMs provide the full potential memory bandwidth. These are the best memory configurations for performance. Balanced memory configurations can also be achieved with two, four, and eight DIMMs, but they do not populate all the memory channels which reduces their memory bandwidth and performance.

Balanced memory configurations are the only memory configurations that should be used if memory bandwidth and performance are important.

Table 1 Summary of all the memory configurations

Number of DIMMs populated	Configuration	Number of Interleave Sets	Relative Performance	Balanced or Unbalanced
1	1:0:0,0:0:0	1	18%	Unbalanced
2	1:0:0,1:0:0	1	35%	Balanced
2	1:1:0,0:0:0	1	35%	Unbalanced
3	1:1:1,0:0:0	1	51%	Unbalanced
3	1:1:0,1:0:0	2	18%	Unbalanced
4	1:1:0,1:1:0	1	66%	Balanced
4	1:1:1,1:0:0	2	35%	Unbalanced
5	1:1:1,1:1:0	2	18%	Unbalanced
6	1:1:1,1:1:1	1	95%	Balanced

Number of DIMMs populated	Configuration	Number of Interleave Sets	Relative Performance	Balanced or Unbalanced
7	2:1:1,1:1:1	2	18%	Unbalanced
7	2:2:0,1:1:1	3	18%	Unbalanced
8	2:2:0,2:2:0	1	69%	Balanced
8	2:1:1,2:1:1	2	35%	Unbalanced
9	2:2:2,1:1:1	2	51%	Unbalanced
9	2:2:1,2:1:1	3	18%	Unbalanced
10	2:2:1,2:2:1	2	69%	Unbalanced
10	2:2:2,2:2:0	2	18%	Unbalanced
11	2:2:2,2:2:1	3	18%	Unbalanced
12	2:2:2,2:2:2	1	100%	Balanced

Near-balanced memory configurations

It is not required that only identical DIMMs be used to achieve good memory bandwidth and performance. A mix of two different capacity DIMMs may be the most cost-effective way to produce the needed memory capacity and performance.

Memory configurations that use only identical DIMMs and follow all the balanced memory guidelines are referred to as *Balanced*. Memory configurations that use two different DIMMs and follow all the balanced memory guidelines are referred to as *Near-Balanced*.

Measurements show about a 3% loss in memory bandwidth when using a Near-Balanced memory configuration as compared to a Balanced memory configuration as long as there are an even number of total memory ranks on each memory channel.

Maximizing memory bandwidth

In order to maximize the memory bandwidth of a server, the following rules should be followed:

- Balance the memory across the processor sockets - all processor sockets on the same physical server should have the same configuration of DIMMs
- Balance the memory across the memory controllers – all memory controllers on a processor socket should have the same configuration of DIMMs
- Balance the memory across the populated memory channels - all populated memory channels should have the same total memory capacity and the same total number of ranks

For optimal memory bandwidth, do the following:

1. Determine your needed memory capacity,
2. Divide this memory capacity by twelve to determine the minimum memory DIMM capacity needed,

3. Round this calculated DIMM capacity up to the closest available memory DIMM capacity, and
4. Populate your server with twelve DIMMs of that capacity.

For example if 256 GB of total memory capacity is needed, each DIMM needs to be a bit more than 21 GB. The closest available memory DIMM size is 32 GB. Therefore, populate your server with twelve 32 GB DIMMs.

A good alternative would be to use a 32 GB and a 16 GB DIMM on each memory channel which provides 6 x (32 GB + 16 GB) = 288 GB of total memory capacity. This alternative is a Near-Balanced memory configuration.

If your system's memory capacity requirement is small, you could divide by six in step 2 above and use only six DIMMs.

You should create a Balanced or Near-Balanced memory configuration by following all the balanced memory guidelines to maximize memory bandwidth and overall server performance.

Summary

Overall server performance is affected by the memory subsystem which can provide both high memory bandwidth and low memory access latency when properly configured. Balancing memory across the memory controllers and the memory channels produces memory configurations that can efficiently interleave memory references among its DIMMs, producing the highest possible memory bandwidth.

An unbalanced memory configuration can reduce the total memory bandwidth to as low as 18% of a balanced memory configuration.

Implementing all three of the balanced memory guidelines described in this brief results in balanced memory configurations producing the best possible memory bandwidth and overall performance.

For more information about Lenovo ThinkSystem servers, go to the following web page:

http://www.lenovo.com/thinksystem

Authors

This paper was produced by the following team of specialists:

Dan Colglazier is a Senior Engineer in the Lenovo Data Center Group Performance Laboratory in Morrisville, NC. He had previously spent over 30 years at IBM, where he started out in the IBM Data Processing Division performance team helping improve the designs of future mainframe storage hierarchies. Assessing the performance of System x® and ThinkSystem servers in both IBM and Lenovo has been his main focus. He is currently the leader of the Performance Design Guidance team. Dan holds a Bachelor of Science degree in Computer Engineering and a Master of Science degree in Electrical Engineering, both from the University of Illinois.

Joe Jakubowski is a Principal Engineer in the Lenovo Data Center Group Performance Laboratory in Morrisville, NC. Previously, he spent 30 years at IBM. He started his career in

the IBM Networking Hardware Division test organization and worked on various Token Ring adapters, switches and test tool development projects. He has spent the past 24 years in the server performance organization focusing primarily on database, virtualization and new technology performance. His current role includes all aspects of x86 server architecture and performance. Joe holds Bachelor of Science degrees in Electrical Engineering and Engineering Operations from North Carolina State University and a Master of Science degree in Telecommunications from Pace University.

Jamal Ayoubi is a Systems Performance Verification Engineer in the Lenovo Data Center Group Performance Laboratory in Morrisville, NC. Before his current position, Jamal interned with the Lenovo Workstation team where he primarily worked on competitive benchmarking and technical solutions. His current role includes CPU, Memory, and PCIe subsystem analysis and performance validation against functional specifications and vendor targets. Jamal holds Bachelor of Science degrees in Electrical Engineering and Computer Engineering from North Carolina State University.

Thanks to the following people for their contributions to this project:

- Deeily Bonieck La Rosa
- David Watts, Lenovo Press

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This document was created or updated on March 26, 2019.

Send us your comments via the **Rate & Provide Feedback** form found at
http://lenovopress.com/lp1089

Trademarks

Lenovo, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at http://www.lenovo.com/legal/copytrade.html.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®	System x®
Lenovo(logo)®	ThinkSystem™

The following terms are trademarks of other companies:

Intel, Intel Xeon, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.